

Causal Video Object Segmentation From Persistence of Occlusions

UCLA-CSD-TR 150002

Brian Taylor Vasiliy Karasev Stefano Soatto

UCLA Vision Lab, University of California, Los Angeles, CA 90095

btay@cs.ucla.edu

vasiliykarasev@ucla.edu

soatto@ucla.edu



Figure 1: Sample outcomes of our scheme: background $c(x) = 0$ (gray) and foreground layers $c(x) = 1$, $c(x) = 2$, $c(x) = 3$ indicated by ■, ■, ■ respectively. On the far right, our algorithm correctly infers that the bag strap is in front of the woman’s arm, which is in front of her trunk, which is in front of the background. Project page: <http://vision.ucla.edu/cvos/>

Abstract

Occlusion relations inform the partition of the image domain into “objects” but are difficult to determine from a single image or short-baseline video. We show how long-term occlusion relations can be robustly inferred from video, and used within a convex optimization framework to segment the image domain into regions. We highlight the challenges in determining these occluder/occluded relations and ensuring regions remain temporally consistent, propose strategies to overcome them, and introduce an efficient numerical scheme to perform the partition directly on the pixel grid, without the need for superpixelization or other preprocessing steps.

1. Introduction

Partitioning the image domain into regions that correspond to “objects” is elusive absent an explicit definition of objects that has a measurable correlate in the image. Gestalt principles [33] provide grouping criteria: continuity, regularity, proximity, compactness, the last of which (figure/ground, or occlusion) is best informed by video. Occlusions have been used extensively for grouping [32, 5, 8, 3]. A feature of [3] is that grouping is obtained via a linear program: local ordering constraints provided by *occluder/occluded relations* are integrated to globally partition the image domain into *depth layers*. The challenge is that errors in determining occlusion relations can have a cascading effect.

Occlusions are usually detected from the residual of optical flow, but even assuming this detection is correct, *occluder relations* are non-trivial to determine. As we show in Fig. 2, correct determination of the occluder requires either knowledge of the motion of the occluded region (which is

undefined), or knowledge of its partition into regions. Hence the conundrum: to determine occlusion relations, so that objects can be segmented, we need to know the objects in the first place. The *first contribution* of our work is to break the conundrum by leveraging motion and appearance priors to hallucinate motion in the occluded region. With the *occluder/occluded* relations we can obtain a depth-layer partition for the image domain. In video, however, nuisances such as motion blur, quantization, scale, and lack of motion can cause layer segmentation to fail. Thus, the *second contribution* is a causal framework for integrating occlusion cues exploiting temporal consistency priors to partition the video into depth layers. Our *third contribution* is to make the solution of the resulting optimization problem efficient using a primal-dual scheme. Our proposed method is competitive to state-of-the-art approaches qualitatively in visual boundaries and quantitatively in numerical benchmarks, while processing video sequences causally, rather than in batch. Samples from our scheme are shown in Fig. 1.

The paper is organized as follows: we set up our problem in Sec. 2. We describe our first contribution in determining occluder relations in Sec. 2.1 and how we leverage prior work [3] in Sec. 2.2. Sec. 3 explores how we causally integrate cues to construct priors for foreground regions in Sec. 3.1, obtain persistent object boundaries in Sec. 3.2, and aggregate occluder relations in Sec. 3.3. Our final model is presented in Sec. 3.4. Implementation and optimization details are covered in Sec. 4–5, including our approach for hallucinating motion in the occluded regions in Sec. 4.2. Empirical evaluation appears in Sec. 6, where we show that the typical failure modes of prior approaches stemming from unreliable occlusion relations are mitigated.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAY 2015		2. REPORT TYPE		3. DATES COVERED 00-00-2015 to 00-00-2015	
4. TITLE AND SUBTITLE Causal Video Object Segmentation From Persistence of Occlusions				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California, Los Angeles (UCLA), UCLA Vision Lab, Los Angeles, CA, 90095				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Occlusion relations inform the partition of the image domain into ???objects??? but are difficult to determine from a single image or short-baseline video. We show how long-term occlusion relations can be robustly inferred from video, and used within a convex optimization framework to segment the image domain into regions. We highlight the challenges in determining these occluder/occluded relations and ensuring regions remain temporally consistent, propose strategies to overcome them, and introduce an efficient numerical scheme to perform the partition directly on the pixel grid, without the need for superpixelization or other preprocessing steps.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

1.1. Related Work

A large number of methods have been proposed for partitioning a video sequence into non-overlapping regions with unique labels, using motion, appearance or their combination [25, 7, 18, 31, 22, 11, 34, 16, 36, 24, 23]. These approaches are susceptible to oversegmentation, which video *object* segmentation attempts to mitigate by assigning a single label to each object. The problem can be cast as multi-label classification, in which a unique label is attached to each object [23], or as binary “foreground”/“background” (FG/BG) classification [12, 22, 36, 24]. While our work produces *depth layers*, and not object labels, these could be added post-mortem.

Many approaches operate *offline* (or *non-causally*), with the entire video available for processing [23, 22, 36, 24], which scales poorly with sequence length, although “streaming” approaches can be used [31, 34]. Our approach is *on-line* (or *causal*), and is closely related to tracking [25, 4, 10], which, unlike us, requires manual initialization.

Estimation of segmentation masks, motion, and depth ordering can be formulated jointly [13, 32, 5, 20, 27, 19, 21, 29, 26, 10, 35], but the resulting problem is nonconvex and requires a substantial computational effort. We separate motion estimation from segmentation and depth ordering, and focus on the latter, which makes a scalable convex formulation possible.

2. Video segmentation with layers

Let $I_t : D \rightarrow \mathbb{R}^3$ be an image of a video $\{I_t\}_{t=1}^T$ defined on the domain $D \subset \mathbb{R}^2$. We seek to partition D into regions, each associated with an integer *depth order*, represented by a function $c_t : D \rightarrow \mathbb{Z}_+$ indicating to which layer each pixel belongs. A layer is then $c_t^{-1}(i) = \{x \in D | c_t(x) = i\}$, where $c_t(x) = 0$ denotes the background and larger values of c_t indicate “foreground” regions $c_t(x) = 1, 2, 3, \dots$. The connected components of non-zero regions correspond to individual objects. It was shown by [5, 3] that depth layers can be inferred from occlusion phenomena, that occur as a result of object or viewer motion, causing parts of the scene to become hidden and others revealed. These inform local order relations between surfaces in the scene: when a surface becomes *occluded*, the image region where it projected to becomes occupied by the *occluder*, which is therefore closer to the viewer. These occluder-occluded relationships can be used as cues for segmenting regions in the image that back-project to distinct objects in the scene.

2.1. The “occluder” and the “occluded”

Under the assumptions of Lambertian reflection, constant illumination, and co-visibility typically implicit in most optical flow algorithms, $I_t(x)$ is related to $I_{t+1}(x)$ by the brightness-constancy equation

$$I_t(x) = I_{t+1}(w_t^{t+1}(x)) + n_t(x), \quad x \in D \setminus \Omega_t^{t+1}(x), \quad (1)$$

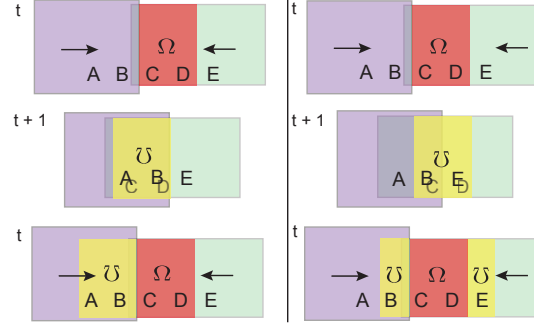


Figure 2: *Initial (top) and final (middle) views of a smaller square sliding under a larger one, producing an occluded region Ω in red (subscripts dropped for readability). Two alternate hypotheses (left and right) for the occluder ($\bar{\Omega}$) in yellow produce different constraints (bottom). Left: Ω moves with E and slides under $A \cup B$. Right: the occluder is split in two—B occludes C and E occludes D. Disambiguation requires either knowledge of the motion in Ω , which is undetermined as it is occluded, or the object segmentation, which is the final goal.*

where w_t^{t+1} is the deformation field that warps the domain of I_t into I_{t+1} and n_t lumps together all un-modeled phenomena and violations of the assumptions. Often, w_t^{t+1} is represented by the optical flow field v_t^{t+1} by $w_t^{t+1}(x) = x + v_t^{t+1}(x)$. The above holds on the entire image domain except in the *occluded regions* Ω_t^{t+1} , where surfaces visible at time t are no longer visible at $t + 1$. In this region, the optical flow is not defined, but can be *extrapolated* from the “co-visible” regions via regularization. Occluded regions are easy to find as a byproduct of optical flow estimation [2], as they yield a large residual n_t via backward flow. What is not easy to find is the *occluder*.

The defining characteristic of the occluder point $y_t^c \in \bar{\Omega}_t^{t+1}$ (the occluder region) corresponding to the occluded point $y_t \in \Omega_t^{t+1}$ (the occluded region) is

$$w_t^{t+1}(y_t^c) = y_t. \quad (2)$$

This equation is somewhat unintuitive as the left hand-side lives in the domain of the image at time $t + 1$ whereas the right-hand side is defined only at time t . This can be interpreted as

$$y_t^c = w_{t+1}^t(y_t), \quad (3)$$

which is completely agnostic of the motion of the occluded region.

Consider Fig. 2: The occluded region, $C \cup D$, could slide under the larger rectangle, and become occluded by $A \cup B$. However, C and D could also actually correspond to different objects, and move independently. In this case, B could be the occluder of C and E could be the occluder of D. To disambiguate between these two hypotheses, we need to know either the motion of D, which is not possible since it

is occluded, or the object partition, which is our goal in the first place. In the example in question, using (2) would favor the hypothesis of B occluding C and E occluding D (right half of Fig. 2). This would yield two ordering constraints, $c(B) > c(C)$ and $c(E) > c(D)$ that hinge on the occluded region and impose no constraints between the visible regions B and E. The latter constraint is also incorrect in the example (Fig. 2 bottom right).

However, while the motion in the occluded region is not determined, it can be hallucinated exploiting regularization priors. Even with a coarse estimate of the motion of D, we could determine if it moves similarly to E, in which case it cannot be occluded by it and must instead be occluded by B. Therefore, in our approach we extrapolate motion to the occluded region, so as to attribute it to a possible occluder. In Sec. 4.2, we discuss how to exploit natural image and motion priors to achieve this. Of course, one could resort to such priors and photometric characteristics of the occluded region to directly determine the grouping of C, D, and E. But again if this was easy, we would have already solved the problem of object segmentation.

2.2. From local ordering constraints to layers

In [3], the following convex model for inferring c_t from occlusion cues was proposed:

$$\begin{aligned} c_t = \arg \min_{c_t: c_t \geq 0} \int_D g_t(x) |\nabla c_t(x)| dx \\ \text{s.t. } c_t(y^c) - c_t(y) \geq 1 \quad \forall (y^c, y) \in O_t. \end{aligned} \quad (4)$$

O_t denotes the set of occlusion cues composed of pairs (y^c, y) , where y^c lies on the occluding surface, and y lies on the surface that was (will be) occluded in the previous (next) frame. The objective $\int_D g_t(x) |\nabla c_t(x)| dx$ is just weighted total variation (TV), with the data-dependent affinity weights (denoted by $g_t(x)$) being small at image and motion boundaries and large otherwise. Note that the “data-term” enters the optimization as a set of constraints which require occluded-occluder pairs to lie in different layers: specifically, the occluder must lie in the layer closer to the viewer (higher values of c_t). An overview of this approach is shown in Fig. 3. While this optimization problem relaxes the integer constraint ($c_t : D \rightarrow \mathbb{Z}_+$), empirically the solutions are piecewise constant and integer valued.

Although this model is formulated for a single time instant t , three frames ($t-1, t, t+1$) are necessary to obtain occlusion cues. However, they are typically not sufficient when small inter-frame motion produces unreliable occlusion constraints. Next, we exploit temporal persistence to overcome this problem.

3. Incorporating motion cues causally

Our causal framework leverages a rich history of image frames, the segmentation cues from those frames (occlu-




Figure 3: Left: The motion of two objects generate occlusions and disocclusions (both denoted by Ω , shown in red). Middle: each occluded region is attributed to a local occluder (\mathcal{U} , shown in yellow). Occluder-occluded relationship constrains objects’ depth-order. Right: resulting depth layers.

sions and weights), and previous layer estimates to facilitate segmentation in the current frame. Large-displacement propagation of these cues via w_t^{t-1} is unstable, rendering cues unusable. But when the motion becomes large, occlusions become easier to detect, making the past unnecessary for segmentation. Thus, these cues are complementary—when the motion is large, sufficient occlusion cues are produced, and w_t^{t-1} is erroneous. When the motion is small, occlusion cues are few, but propagation is reliable. This motivates an adaptive integration of cues based on motion. A weighting function $m_t(x) \doteq \alpha \exp(-|v_t^{t-1}|/\mu_v)$ is used, where $\alpha \in [0, 1]$, v_t^{t-1} is the optical flow, and μ_v is the mean value of v for this frame. The weight decreases with large motion, regardless of how long ago it occurred. The following sections describe the temporal cues leveraged in our framework. Note that the variable being optimized over is always c_t , and c_{t-1} is always available as a result of previous optimization.

3.1. Once an object, always an object

Layer values $c_t(x)$ are not constant over time, as objects can move in front of one another and switch order of their distance to the viewer. However, once an object is detected, it should not later be labeled as background—even if it stops moving and produces no occlusion cues for segmentation.

This can be enforced causally using the prior segmentation result (c_{t-1}) via a (convex) constraint:

$$c_t(x) \geq 1 \quad \forall x \in F, \quad F = \{c_{t-1}(w_t^{t-1}(x)) \geq 1\} \quad (5)$$

where F is the indicator of the previous frame’s foreground region warped into the current frame. To mitigate errors in prior segmentations, we relax the constraint and penalize violations with a hinge loss:

$$\int_D \kappa_t(x) \max(0, 1 - c_t(x)) dx \quad (6)$$

with κ_t being the cost of violating the constraint. Choosing $\kappa_t(x) = 0$ for x outside F allows us to write the penalty as an integral over entire image domain D . As $\kappa_t(x) \rightarrow \infty$ for $x \in F$, the hard constraint (5) is recovered.



Figure 4: c_{t-1} (column 1) is used to compute the foreground prior (κ_t) (column 2). Without κ_t , the resulting c_t completely misses the objects (column 3), however with κ_t , c_t succeeds (last column). Note c_{t-1} and c_t look very similar— κ_t helps most during small-baseline motion when occlusion cues are weak but c_{t-1} easily predicts c_t .

The cost of violating the constraint is computed recursively, with initial condition $\kappa_1(x) = 0$, as

$$\kappa_t(x) = m_t(x)\kappa_{t-1}(w_t^{t-1}(x)) + \mathbb{1}\{c_{t-1}(w_t^{t-1}(x)) \geq 1\}$$

where $\mathbb{1}$ is a characteristic function ($\mathbb{1}\{X\} = 1$ if X is true, and is 0 otherwise). This *foreground prior* boosts $\kappa_t(x)$ wherever the corresponding points are labeled as foreground in the previous frame and diminishes it over time and motion as described above. As demonstrated in Fig. 4, whenever motion is small, instantaneous occlusion cues are insufficient to perform segmentation, and this notion of temporal consistency is helpful.

To avoid the entire image domain from becoming foreground, we introduce an additional regularization penalizing layer values

$$\tau \int_D c_t(x) dx. \quad (7)$$

This is similar to the regularization used in [3], although they use the ℓ_∞ norm, whereas here we use ℓ_1 . This term encourages pixels to lie in the background layer, unless sufficient evidence pushes them into the foreground.

3.2. Persistent layer boundaries

While depth-layer *values* are not persistent, their boundaries are. Unless objects split or merge, we have

$$\mathbb{1}\{\nabla c_t(x) \neq 0\} = \mathbb{1}\{\nabla c_{t-1}(w_t^{t-1}(x)) \neq 0\}. \quad (8)$$

This is a nonconvex constraint. However, enforcing $\nabla c_t(x) = 0$ wherever $c_{t-1}(w_t^{t-1}(x)) = 0$ is simple (a linear constraint), and its relaxed version with a hinge loss and associated cost $u_t(x)$ is equivalent to increasing weights in TV regularization (shown in appendix). This leaves the hard part: enforcing $\nabla c_t(x) \neq 0$ wherever $c_{t-1}(w_t^{t-1}(x)) \neq 0$. To remain within a convex optimization framework, we treat this as a bias and set the corresponding $u_t(x)$ to be negative, which *decreases* the corresponding TV weights (which are kept nonnegative to preserve convexity). This *layer unity*




Figure 5: Occlusion cues from the current frame alone (O_t), with occluded points (Ω) in red and occluder points (\mathcal{U}) in yellow, (column 1) fail to segment the objects (column 2). However, aggregating constraints over time (\bar{O}_t) (column 3) successfully recovers all of them (last column).

prior is also computed recursively, with $u_1(x) = 0$, as

$$u_t(x) = m_t(x)u_{t-1}(w_t^{t-1}(x)) + \mathbb{1}\{\nabla c_{t-1}(w_t^{t-1}(x)) = 0\} - \mathbb{1}\{\nabla c_{t-1}(w_t^{t-1}(x)) \neq 0\}.$$

We also perform temporal aggregation of the TV affinity weights. In each frame, we compute the *boundary strength* $\rho_t(x) \in \mathbb{R}_+$, as described in Sec. 4. The aggregated boundary strength $\bar{\rho}_t(x)$ is (with $\bar{\rho}_1(x) = 0$)

$$\bar{\rho}_t(x) = m_t(x)\bar{\rho}_{t-1}(w_t^{t-1}(x)) + \rho_t(x). \quad (9)$$

The aggregated TV weights used in the optimization are

$$g_t(x) = \max(0, 1 - \bar{\rho}_t(x) + u_t(x)). \quad (10)$$

3.3. Occlusion cue aggregation

Instantaneous occlusion constraints (O_t) are accumulated into the aggregated constraints set $\bar{O}_t = w_{t-1}^t(\bar{O}_{t-1}) \cup O_t$, where past constraints \bar{O}_{t-1} are propagated to the current frame by the motion of the occluder $w_{t-1}^t(y^c)$ (see Fig. 5). The base condition is $\bar{O}_1 = O_1$. The constraint penalty weights λ , computed by (4.1), are adjusted over time by

$$\lambda_{t,i} = m_t(y_i^c)\lambda_{t-1,i} + \mathbb{1}\{c_{t-1}(w_t^{t-1}(y_i^c)) \geq c_{t-1}(w_t^{t-1}(y_i))\}.$$

3.4. Overall model

The final model that incorporates occlusion cues, weights, foreground and unity priors is

$$\begin{aligned} c_t = \arg \min_{c_t \geq 0} & \int_D g_t(x) |\nabla c_t(x)| dx + \tau \int_D c_t(x) dx \\ & + \int_D \kappa_t(x) \max(0, 1 - c_t(x)) dx \\ & + \sum_{i=1}^N \lambda_i \max(0, 1 - c_t(y_i^c) - c_t(y_i)), \end{aligned} \quad (11)$$

where the first term (weighted TV) ensures that the result is piecewise constant, the second term (foreground prior) encourages regions to have nonzero layer values wherever $\bar{\kappa}_t(x)$ is large, the third (model selection) term prevents the creation of spurious layers, and the fourth is the penalty for violating the occlusion constraints.

4. Implementation details

For each frame, we incorporate appearance, edge, and motion information into the weights $\rho_t(x)$ in (9) as follows:

$$\rho_t(x) = 1 - (\beta_I h(|\nabla I(x)|) + \beta_E h(E(x)) + \beta_v h(|\nabla v_t^{t+1}(x)|))$$

where $h(x) = \exp(-x/\mu_x)$, μ_x is the average value of x . $E(x) \in [0, 1]$ is the output of an edge detector [14] with $E(x) \approx 1$ at the boundaries. In our experiments, $(\beta_I, \beta_E, \beta_v) = (0.2, 0.4, 0.4)$. Following [24], we also adjust the motion term by the difference in flow angles at the pixels where flow magnitude is small.

4.1. Occlusion constraint weights

Often the occluded and occluding surfaces differ in appearance, motion, and are separated by a strong image boundary, suggesting λ be computed in a fashion similar to (4):

$$\lambda_i = \eta_i (1 - (\beta_I h(|I(y_i^c) - I(y_i)|) + \beta_E h(\hat{E}(y_i^c, y_i)) + \beta_v h(|v_t^{t+1}(y_i^c) - v_t^{t+1}(y_i)|)))$$

where the gradient operator is replaced by a difference between appearance, edge, and motion statistics of y^c and y . Here, $E(x)$ is replaced by $\hat{E}(x_1, x_2)$ —the strongest edge response on the line connecting y^c and y . We additionally validate y^c and y as an occluder-occluded pair with weight η , which measures the degree to which y and y^c move toward each other. Indeed, unless they do so, Ω_t^{t+1} cannot take the place of Ω_t^{t+1} , i.e. when $\eta(y^c, y)$ in

$$\Delta(y^c, y) = (v_t^{t+1}(y^c) - v_t^{t+1}(y))^T \left(\frac{y^c - y}{\|y^c - y\|} \right) \quad (12)$$

$$\eta(y^c, y) = \max(0, 1 - \exp(-\theta \Delta(y^c, y)))$$

is small, then y^c is less likely to occlude y . We choose $\theta = 2$ so that $\Delta(y^c, y) = 1$ yields a high score. $\lambda_i \approx 1$ whenever the appearance and motion of y^c and y are “different” and the points are moving toward each other. Finally, assuming that the occluded and occluding surfaces differ in appearance, we can locally perturb constraints with the goal of correcting them; this procedure is described in the appendix. Altogether, these factors alter the constraints to help us discount potentially erroneous cues, which occur due to inevitable errors in optical flow and occlusion estimation.

4.2. Flow extrapolation over the occlusion region

As noted in Sec. 2.1, $v_t^{t+1}(x)$ for $x \in \Omega_t^{t+1}$ is undefined (1) and filled in by the regularizer, which corresponds to enforcing priors on motion. The simplest priors rely solely on continuity, tending to smooth motion boundaries, while more sophisticated ones attempt to preserve them. We use the cross-bilateral filter [15] to enforce such priors on v_t^{t+1}




Figure 6: Cross bilateral filtering extrapolates flow in Ω via motion and appearance priors, facilitating reliable occluder determination. Left: The extrapolated motion field \hat{v}_t^{t+1} . Boxes highlight occluded regions where notable change (often improvement) occurs. Right: For each box, v_t^{t+1} (top) and \hat{v}_t^{t+1} (bottom) are shown.

in the occluded regions based on the backward flow v_t^{t-1} :

$$\hat{v}_t^{t+1}(z) = \frac{1}{V_z} \int_D v_t^{t+1}(x) P(x \notin \Omega_t^{t+1}) \mathcal{G}(v_t^{t-1}(x) - v_t^{t-1}(z), \sigma_v) \mathcal{G}(x - z, \sigma_x) dx, \quad (13)$$

where \hat{v}_t^{t+1} is the extrapolated forward flow, $P(x \notin \Omega_t^{t+1})$ is the probability of x being visible, \mathcal{G} is the gaussian kernel $\mathcal{G}(x, \sigma) = \exp(-\|x\|^2/2\sigma^2)$, and V_z is a normalization term. We can filter the backward flow \hat{v}_t^{t-1} by exchanging $t + 1$ with $t - 1$ and vice versa. Extrapolating flow is key to determining the occluder (Fig. 2), but cannot be proven “correct” as it hinges critically on the choice of prior. v_t^{t+1} is computed using publicly-available code [28] (“classic-nl”), and occlusions are computed by thresholding the residual image. See appendix for further details.

4.3. Foreground prior region

In practice, motion estimation makes mistakes near object boundaries (e.g. occluded regions). When computing κ_t , we first warp κ_{t-1} to the current frame and then use morphological operations to erode the edges proportionally to the magnitude of the flow in that region. This ensures the prior does not leak outside of the object regions, but produces a poor estimate near the boundaries. To help recover the structure of these edges, we incorporate a set of local shape classifiers as in [4] to better capture and predict the shape of the object boundary, the details of which are in the appendix.

5. Optimization

The optimization problem (11) is convex but large enough that off-the-shelf methods cannot solve it without resorting to superpixels or other pre-processing to reduce its dimension. Here we present an efficient numerical primal-dual scheme based on [9] that allows us to solve it on the pixel grid.

The indicator function – not to be confused with characteristic function $\mathbb{1}$ used above – of a set A is defined by $I_A(x) = 0$ for $x \in A$ and $I_A(x) = \infty$ for $x \notin A$. For a function f , the *convex conjugate* is defined as $f^*(y) \doteq$

Algorithm 1 Layer Solver

Initialize: Pick $\sigma_y, \sigma_c > 0$, $\sigma_y \sigma_c \leq \frac{1}{8}$, and $\theta \in [0, 1]$. Arbitrarily initialize feasible y_1^0, y_2^0, c^0 . Set $\bar{x}^0 = c^0$.

Perform iterates for $k = 0, 1, 2, \dots$:

$$\begin{aligned} y_1^{k+1} &= \text{prox}_{\sigma_y F_1^*}(y_1^k + \sigma_y \mathcal{D} \bar{x}^k) \\ y_2^{k+1} &= \text{prox}_{\sigma_y F_2^*}(y_2^k + \sigma_y \mathcal{D}_{occ} \bar{x}^k) \\ c^{k+1} &= \text{prox}_{\sigma_c G}(c^k - \sigma_c (\mathcal{D}^T y_1^{k+1} + \mathcal{D}_{occ}^T y_2^{k+1})) \\ \bar{x}^{k+1} &= c^{k+1} + \theta(c^{k+1} - c^k). \end{aligned}$$

$\sup_x y^T x - f(x)$. The *prox operator* of f is defined as

$$\text{prox}_{\sigma f}(y) \doteq \arg \min_x \frac{1}{2\sigma} \|x - y\|^2 + f(x). \quad (14)$$

Since the optimization is performed on a finite pixel grid, the depth values c can be written as a vector $c \in \mathbb{R}_+^n$, with c_i indicating the layer value at the i -th pixel. We denote by \mathcal{D} the gradient operator represented by a matrix of finite differences. Weights associated with the edges are denoted by the diagonal matrix W . A difference matrix for occlusion constraints is denoted by \mathcal{D}_{occ} and the cost of violating constraints by λ . As before, τ is used for regularization and κ is a weighted indicator of the foreground region. We can then write the objective in shorthand as

$$\begin{aligned} \min_c \|W\mathcal{D}c\|_1 + \tau^T c + \kappa^T \max(0, 1 - c) + \\ \lambda^T \max(0, 1 - \mathcal{D}_{occ}c) + I_{\{c \geq 0\}}(c). \end{aligned} \quad (15)$$

Let $G(c) = \tau^T c + \kappa^T \max(0, 1 - c) + I_{\{c \geq 0\}}(c)$. Also, let $z_1 = \mathcal{D}c$, $z_2 = \mathcal{D}_{occ}c$, construct the functions $F_1(z_1) = \|Wz_1\|_1$, $F_2(z_2) = \lambda^T \max(0, 1 - z_2)$, and introduce the dual variables y_1, y_2 . The augmented Lagrangian follows as

$$\begin{aligned} \min_{z_1, z_2, c} \max_{y_1, y_2} F_1(z_1) + F_2(z_2) + G(c) + \\ y_1^T (\mathcal{D}c - z_1) + y_2^T (\mathcal{D}_{occ}c - z_2), \end{aligned} \quad (16)$$

or, equivalently, using the convex conjugates, as

$$\min_c \max_{y_1, y_2} G(c) - F_1^*(y_1) - F_2^*(y_2) + y_1^T \mathcal{D}c + y_2^T \mathcal{D}_{occ}c.$$

This saddle-point problem is addressed in [9], so we can apply their primal-dual algorithm shown in Alg. 1.

Alg. 1 depends on the ability to compute proximal operators for G , F_1^* and F_2^* . All three operators have simple closed form solutions that require few arithmetic operations:

$$\text{prox}_{\sigma G}(y) = \max(0, \min(y - \sigma\tau + \sigma\kappa, \max(1, y - \sigma\tau))) \quad (17)$$

$$\text{prox}_{\sigma F_1^*}(y) = \text{sign}(y) \min\{\text{diag}(W), |y|\} \quad (18)$$

$$\text{prox}_{\sigma F_2^*}(y) = \min(\max(y - \sigma 1, -\lambda), 0) \quad (19)$$

Derivation details are reported in the appendix.

6. Experiments

Our method segments video into depth layers. Unfortunately, no benchmark dataset is available to evaluate it directly. However, our method can be modified to produce binary and multi-label segmentations; leveraging this, we evaluate the algorithm on two datasets: MoSeg [23] (designed for video *object* segmentation with no consideration for depth ordering), on which we focus, as well as BVSD [17] (designed for video segmentation).

Evaluation methodology. We follow the process described in [23]. The dataset contains 59 sequences, ranging from 19 to 800 frames. Each has pixel-wise ground truth annotation for a sparse subset of frames (3–41). As in [23], we report *precision*, *recall*, *F-measure*, and the number of extracted objects (regions with F-measure ≥ 0.75). For multi-label segmentation tasks, we treat each connected component of the depth layers as a unique “object”. We also evaluate on foreground/background (FG/BG) video object segmentation, which come directly from depth layers as $FG \doteq \{x : c(x) \geq 1\}$, $BG \doteq \{x : c(x) = 0\}$. *Precision*, *recall*, and *F-measure* are reported on the ground truth annotations converted to binary masks. Note we cannot evaluate “number of extracted objects” in the FG/BG scenario.

The methods we compare against ([18, 22, 24, 23]) are non-causal and “batch”, whereas our method is causal. Since we do not know the future, we do not detect objects until they undergo sufficient motion, which sometimes causes us to miss objects in the beginning of video sequences. To fairly compare against non-causal methods, we also perform a non-causal evaluation (reported as “NC”)—we run our algorithm forward in time to accumulate all priors, and then backward in time. The latter half is used for evaluation.

Effects of system components. In Sec. 3 we described individual components of our model and showed examples where they improved results (see Fig. 4, 5). Here we quantify this improvement. We evaluate [3] (“BASIC”), their temporal extension (“TE”), foreground-background prior (Sec. 3.1, “FG”), and the full model (“FULL”). In addition, we evaluated the full model without flow extrapolation (Sec. 4.2) to understand its effects (“NOFE”). These results are reported in Table 1. “BASIC” does not use long-term temporal information. “TE” integrates weights using previous segmentations, increasing the cost of making a cut away from object boundaries. “FG” discourages previously segmented regions from falling into background. “FULL” is a combination of all components.

The “BASIC” method does not use temporal information, so on the multi-label benchmark, whenever objects disappear (as they often do, due to insufficient motion) and re-appear, they are assigned a *new* object label. Long-term integration helps avoid missed detections and propagates object labels throughout the sequence. Performance on the FG/BG evaluation suggests that objects are often not detected at all.

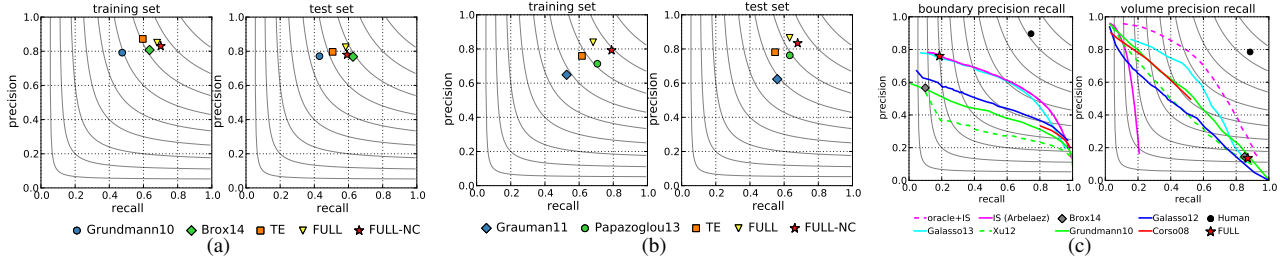


Figure 7: (a-b) Comparison on MoSeg: (a) multi-label segmentation, (b) FG/BG segmentation. (c) Comparison on BVSD.

Multi-label segmentation								
	Training set (29 sequences)				Test set (30 sequences)			
	P	R	F	N/65	P	R	F	N/69
BASIC	84.90	53.10	65.34	10	78.80	44.49	56.87	4
TE	87.20	59.60	70.81	17	79.64	50.73	61.98	7
FG	86.98	60.99	71.71	18	79.04	52.08	62.79	10
NOFE	86.67	58.06	69.54	14	80.71	50.64	62.24	8
FULL	85.00	67.99	75.55	21	82.37	58.37	68.32	17
FULL-NC	83.00	70.10	76.01	23	77.94	59.14	67.25	15
[18]	79.17	47.55	59.42	4	77.11	42.99	55.20	5
[23]	81.50	63.23	71.21	16	74.91	60.14	66.72	20

Binary segmentation								
	Training set (29 sequences)				Test set (30 sequences)			
	P	R	F	-	P	R	F	-
BASIC	89.99	40.86	56.21	-	93.21	33.69	49.49	-
TE	75.94	61.64	68.05	-	78.11	54.68	64.33	-
FG	75.93	63.07	68.91	-	76.97	56.16	64.94	-
NOFE	68.92	66.09	67.48	-	74.27	53.99	62.52	-
FULL	83.92	68.19	75.24	-	86.54	63.20	73.05	-
FULL-NC	79.26	78.99	79.12	-	83.41	67.91	74.87	-
[22]	64.86	52.70	58.15	-	62.32	55.97	58.97	-
[24]	71.34	70.66	71.00	-	76.29	63.29	69.18	-

Table 1: Comparison of our approach (rows 4–5) to baselines using individual components (rows 1–3) and state-of-the-art (rows 6–7) on the MoSeg dataset. R = recall, P = precision, F = F-measure, N = number of extracted objects.

Precision decreases for the “FULL” system due to an increased number of “false positives”—often we detect more objects than labeled in the annotation (see Fig. 8). “NC” provides a small performance boost by allowing us to label objects before they move.

Video object segmentation. In Table 1 and Fig. 7 we report results of the comparison with multi-label dense motion segmentation [23], video over-segmentation [18], as well as binary (i.e. FG/BG) video object segmentation methods [22, 24]. On multi-label segmentation, we outperform [18], [3], and [23] in F-measure. The improvement from the latter is not great; however, note that unlike theirs, our method is causal and has a small memory footprint. We are not the best in terms of “number of extracted objects”. As mentioned before, unless the object undergoes sufficient motion, it will not be detected. On FG/BG segmentation, we outperform [22], [24], and [3].

Video segmentation. BVSD [30, 17] contains 40 training and 60 testing sequences, each up to 121 frames. Pixel-wise ground truth annotation is provided for a subset of frames. Video sequences are in HD; we resize images to 540×960 . While we report results for a variety of algorithms

[11, 1, 16, 34] (with data from [17]), our primary point of comparison is [23]. Performance is benchmarked using “boundary precision-recall” (BPR) and “volume precision-recall” (VPR) metrics. BPR is commonly used in image segmentation, while VPR quantifies the spatiotemporal overlap between machine-generated and ground-truth segmentations (see [17] for details).

Video *object* segmentation algorithms are expected to be in the high-precision regime in BPR, and in the high-recall regime in VPR, which indeed both we and [23] satisfy (see Fig. 7). We obtain $(P, R, F) = (0.760, 0.186, 0.299)$ and $(0.136, 0.870, 0.234)$ on BPR and VPR respectively, while they obtain $(0.566, 0.100, 0.170)$ and $(0.146, 0.852, 0.249)$. Sample results are in Fig. 9. Note that the ground truth is often fine-grained—with objects spanned by multiple regions. Thus, on this benchmark, object segmentation methods will not obtain the best F-measure.

Timing. Given optical flow (which video segmentation often requires as input), our algorithm takes 30s for VGA images on a standard desktop; most of the time is spent solving (11), but a GPU implementation can reduce this.

7. Discussion

Occlusion relations inform the partition of the image domain into segments, but proper inference of such relations requires knowledge of the segments in turn. Rather than tackling an intractable chicken-and-egg problem, we use priors informed by Gestalt principles to arrive at a convex optimization scheme that can be efficiently solved with primal-dual methods. To compare with existing benchmarks, we converted our layers into “objects” and into “foreground/background”. The evaluation highlights strengths and limitations of our method, with some of the latter due to the particular characteristics of the benchmarks. While our scheme still relies on decent optical flow and occlusion detection to bootstrap layer segmentation, it is less prone to cascading failure than previous methods, as it better exploits priors on motion, appearance, and layer consistency.

Acknowledgements. We would like to thank Virginia Estellers for many useful suggestions and discussions on optimization. This work was supported by NSF RI-1422669, ONR N00014-13-1-034, FA8650-11-1-7156.




Figure 8: Sample results on MoSeg. Left to right: original image, ground truth, [22], [24], [18], [23], ours (object maps). Top two rows: occlusion cues allow us to obtain even the barely visible cars (row 1 - orange, row 2 - red, row 2 - green). Row 3: use of both motion and appearance cues allows us to generate an accurate object boundary. Row 4: occlusion cues yield three depth layers (bicyclist, tree, background) (see also Fig. 1). Notice that the tree (and some cars in rows 1–2) is not annotated, so our scheme is penalized despite providing the correct answer. [22, 24] suffer from trailing and only produce binary segmentation. [18] suffers from oversegmentation. [23] performs comparably to our method; The last two rows show failure cases. Row 5: the painting is recognized as an “object” due to false occlusion detection; the hand is assigned to a separate layer. Row 6: the lioness is missed due to insufficient motion and lack of occlusions.



Figure 9: Sample results on BVSD. Left to right: original image, ground truth, [18], [23], ours (object maps). Row 1: as reflected by BPR, our method produces accurate boundaries (see Fig. 7). Both actors are correctly segmented—the arm occluding the animal’s body is a distinct depth layer. Row 2: “failure case”—complex motion and inaccurate flow can result in inaccurate segmentations. Row 3: failure case—object is not detected throughout the sequence due to lack of occlusions.

References

- [1] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *TPAMI*, 33(5), 2011.
- [2] A. Ayvaci, M. Raptis, and S. Soatto. Sparse occlusion detection with optical flow. *International Journal of Computer Vision*, 97(3), May 2012.
- [3] A. Ayvaci and S. Soatto. Detachable object detection: Segmentation and depth ordering from short-baseline video. *PAMI*, 34(10), 2012.
- [4] X. Bai, J. Wang, D. Simons, and G. Sapiro. Video snapcut: robust video object cutout using localized classifiers. In *ACM Transactions on Graphics (TOG)*, 2009.
- [5] L. Bergen and F. Meyer. Motion segmentation and depth ordering based on morphological segmentation. In *ECCV*, 1998.
- [6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [7] W. Brendel and S. Todorovic. Video object segmentation by tracking regions. In *ICCV*, 2009.
- [8] G. J. Brostow and I. Essa. Motion based decompositing of video. In *ICCV*, 1999.
- [9] A. Chambolle and T. Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1), 2011.
- [10] J. Chang and J. W. Fisher III. Topology-constrained layered tracking with latent flow. In *ICCV*, 2013.
- [11] J. J. Corso, E. Sharon, S. Dube, S. El-Saden, U. Sinha, and A. Yuille. Efficient multilevel brain tumor segmentation with integrated bayesian model classification. *IEEE Transactions on Medical Imaging*, 27(5), 2008.
- [12] A. Criminisi, G. Cross, A. Blake, and V. Kolmogorov. Bilayer segmentation of live video. In *CVPR*, 2006.
- [13] T. Darrell and A. Pentland. Robust estimation of a multi-layered motion representation. In *IEEE Workshop on Visual Motion*, 1991.
- [14] P. Dollár and C. L. Zitnick. Structured forests for fast edge detection. In *ICCV*, 2013.
- [15] E. Eisemann and F. Durand. Flash photography enhancement via intrinsic relighting. *ACM Trans. Graph.*, 23(3), 2004.
- [16] F. Galasso, R. Cipolla, and B. Schiele. Video segmentation with superpixels. In *ACCV*, 2012.
- [17] F. Galasso, S. Naveen, T. J. Cardenas, T. Brox, and B. Schiele. A unified video segmentation benchmark: Annotation, metrics and analysis. In *ICCV*, 2013.
- [18] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph based video segmentation. In *CVPR*, 2010.
- [19] J. Jackson, A. J. Yezzi, and S. Soatto. Dynamic shape and appearance modeling via moving and deforming layers. *IJCV*, 2008.
- [20] N. Jojic and B. J. Frey. Learning flexible sprites in video layers. In *CVPR*, 2001.
- [21] M. P. Kumar, P. H. Torr, and A. Zisserman. Learning layered motion segmentations of video. *IJCV*, 76(3), 2008.
- [22] Y. J. Lee, J. Kim, and K. Grauman. Key-segments for video object segmentation. In *ICCV*, 2011.
- [23] P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *TPAMI*, 36(6), 2014.
- [24] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *ICCV*, 2013.
- [25] X. Ren and J. Malik. Tracking as repeated figure/ground segmentation. In *CVPR*, 2007.
- [26] T. Schoenemann and D. Cremers. A coding-cost framework for super-resolution motion layer decomposition. *TIP*, 2012.
- [27] P. Smith, T. Drummond, and R. Cipolla. Layered motion segmentation and depth ordering by tracking edges. *TPAMI*, 26(4), 2004.
- [28] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *CVPR*, 2010.
- [29] D. Sun, E. Sudderth, and M. Black. Layered segmentation and optical flow estimation over time. In *CVPR*, 2012.
- [30] P. Sundberg, T. Brox, M. Maire, P. Arbelaez, and J. Malik. Occlusion boundary detection and figure/ground assignment from optical flow. In *CVPR*, 2011.
- [31] A. Vazquez-Reina, S. Avidan, H. Pfister, and E. Miller. Multiple hypothesis video segmentation from super-pixel flows. In *ECCV*, 2010.
- [32] J. Y. Wang and E. H. Adelson. Representing moving images with layers. *TIP*, 3(5), 1994.
- [33] M. Wertheimer. *Laws of organization in perceptual forms*. W. D. Ellis, 1939.
- [34] C. Xu, C. Xiong, and J. J. Corso. Streaming hierarchical video segmentation. In *ECCV*, 2012.
- [35] Y. Yang and G. Sundaramoorthi. Modeling self-occlusions in dynamic shape and appearance tracking. In *ICCV*, 2013.
- [36] D. Zhang, O. Javed, and M. Shah. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *CVPR*, 2013.

A. Additional results

In the main paper, we reported results on the BVSD dataset. The dataset contains three “subtasks”—“motion”, “camera motion”, and “non-rigid motion” (details in [17]). Results on these subtasks are reported below in Fig. 10 and Table 2.

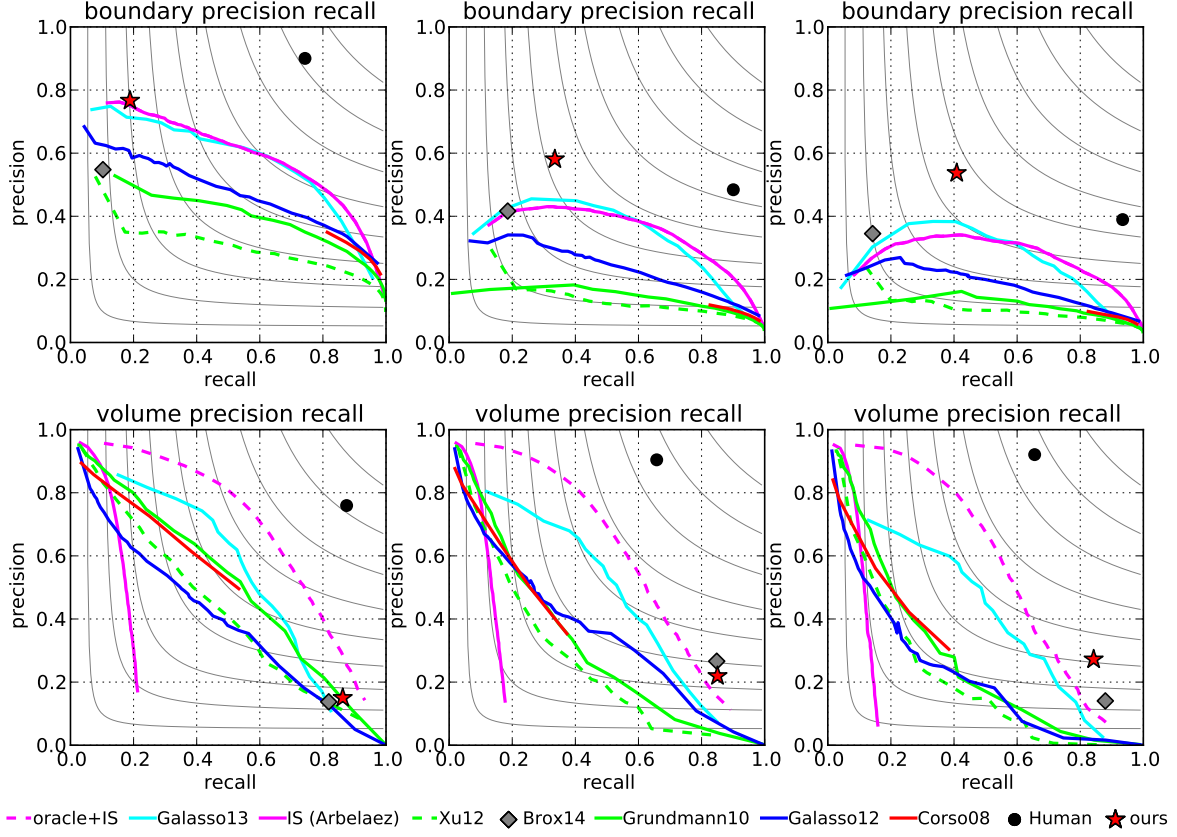


Figure 10: Results on BVSD subtasks; from left to right: “motion”, “camera motion”, “nonrigid motion”

	BPR			VPR		
	P	R	F	P	R	F
Motion	58.07	33.50	42.48	21.93	85.04	34.87
Camera motion	76.65	18.87	30.28	14.99	86.28	25.55
Nonrigid motion	53.71	40.86	46.42	27.20	84.45	41.12
General	76.02	18.65	29.95	13.56	87.09	23.46

Table 2: Precision-recall statistics of our method on BVSD subtasks. “General” subtasks contains all sequences and was reported in main paper.

B. Optimization details

In this section we derive the prox-operators used by the primal-dual algorithm.

B.1. Proximal operator for $G(c) = \tau^T c + \kappa^T \max(0, 1 - c) + I_{\{c \geq 0\}}(c)$

First define $g_i(c_i) = \tau_i c_i + \kappa_i \max(0, 1 - c_i) + I_{\{c_i \geq 0\}}(c_i)$, so that $G(c) = \sum_{i=1}^n g_i(c_i)$. This is done in anticipation of the prox-operator for G being separable¹:

$$\mathbf{prox}_{\sigma G}(y) = \arg \min_{c \geq 0} \frac{1}{2\sigma} \|c - y\|_2^2 + \tau^T c + \kappa^T \max(0, 1 - c) \quad (20)$$

$$= \arg \min_{c \geq 0} \sum_{i=1}^n \frac{1}{2\sigma} (c_i - y_i)^2 + \tau_i c_i + \kappa_i \max(0, 1 - c_i) \quad (21)$$

$$= (\mathbf{prox}_{\sigma g_1}(y_1), \dots, \mathbf{prox}_{\sigma g_n}(y_n)) \quad (22)$$

We now will derive the expression for $\mathbf{prox}_{\sigma g_i}(y_i)$. Since the general form of the expression will be the same for each i , we drop the index for clarity. For convenience, let $f(c) = \frac{1}{2\sigma} (c - y)^2 + g(c)$ (as shown in Fig. 11, f is a parabola with a kink), i.e. $\mathbf{prox}_{\sigma g}(y) = \arg \min f(c)$. Since f is strictly convex and the minimizer is unique, we initially ignore the constraint $\{c \geq 0\}$, eventually projecting the minimizer back onto the feasible domain if necessary. The subdifferential of $f(c)$ is

$$\partial f(c) = \begin{cases} \frac{1}{\sigma}(c - y) + \tau & \text{if } c > 1 \\ \frac{1}{\sigma}(c - y) + \tau - \kappa & \text{if } 0 \leq c < 1 \\ \frac{1}{\sigma}(c - y) + \tau - \kappa[0, 1] & \text{if } c = 1 \end{cases} \quad (23)$$

The optimality condition $0 \in \partial f(c)$ is satisfied when

$$c^* = \begin{cases} y - \sigma\tau & \text{if } y > 1 + \sigma\tau \\ 1 & \text{if } y \in [1 + \sigma(\tau - \kappa), 1 + \sigma\tau] \\ \max(0, y - \sigma(\tau - \kappa)) & \text{if } y < 1 + \sigma(\tau - \kappa) \end{cases} \quad (24)$$

These three cases are shown in Fig. 11,




Figure 11: Top: $f(c)$ shown for the three cases described in (24). Bottom: subdifferentials of corresponding $f(c)$. Shown in red are the pairs $(c^*, f(c^*))$.

which can be rewritten as

$$\mathbf{prox}_{\sigma g_i}(y_i) = \max \left(0, \min \left(y_i - \sigma(\tau_i - \kappa_i), \max(1, y_i - \sigma\tau_i) \right) \right). \quad (25)$$

Using (22) and interpreting max/min componentwise, the prox-operator for G is written as

$$\mathbf{prox}_{\sigma G}(y) = \max \left(0, \min \left(y - \sigma(\tau - \kappa), \max(1, y - \sigma\tau) \right) \right). \quad (26)$$

¹We use the following “separable sum” rule. If $f(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}) = g(x_1) + h(x_2)$, then $\mathbf{prox}_f(\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}) = \begin{pmatrix} \mathbf{prox}_g(x_1) \\ \mathbf{prox}_h(x_2) \end{pmatrix}$

B.2. Proximal operator for $F_1^*(y)$, with $F_1(y) = \|Wy\|_1$

We note that this operator appears in all problems that involve TV-regularization, that our derivation below is by no means novel, and that it is included only for completeness. Since the function is a (weighted) norm, we expect its convex conjugate (i.e. F_1^*) to be a (weighted) indicator of a dual norm ball. Moreover, we expect the proximal operator of the indicator function to be the projection onto the feasible set [6]. The steps below verify it.

$$F_1^*(z) = \sup_y y^T z - F_1(y) \quad (27)$$

$$= \sum_{i=1}^n \max_{y_i} \underbrace{y_i z_i - w_i |y_i|}_{g_i(y_i)} \quad (28)$$

The maximum of each term $g_i(y_i)$ can be determined as

$$\max_{y_i} g_i(y_i) = w_i \left(\max_{y_i} y_i \frac{z_i}{w_i} - |y_i| \right) \quad (29)$$

$$= \begin{cases} 0 & \text{if } \left| \frac{z_i}{w_i} \right| \leq 1 \\ +\infty & \text{else} \end{cases}, \quad (30)$$

so $F_1^*(z)$ is the indicator of the weighed ℓ_∞ ball:

$$F_1^*(z) = \begin{cases} 0 & \text{if } \|W^{-1}z\|_\infty \leq 1 \\ +\infty & \text{else} \end{cases}. \quad (31)$$

Using $\text{diag}(W)$ to write the main diagonal as a vector (i.e. $\text{diag}(W) = (W_{11}, \dots, W_{nn})$), the prox operator for F_1^* can be written as

$$\text{prox}_{\sigma F_1^*}(y) = \arg \min_z \frac{1}{2\sigma} \|z - y\|_2^2 + F_1^*(z) \quad (32)$$

$$= \text{sign}(y) \min\{\text{diag}(W), |y|\}. \quad (33)$$

B.3. Proximal operator for $F_2^*(y)$, with $F_2(y) = \lambda^T \max(0, 1 - y)$

The convex conjugate is

$$F_2^*(z) = \max_y z^T y - F_2(y) \quad (34)$$

$$= \sum_{i=1}^n \max_{y_i} \underbrace{y_i z_i - \lambda_i \max(0, 1 - y_i)}_{g_i(y_i)}. \quad (35)$$

Notice that when $z_i > 0$, $\max_{y_i} g_i(y_i) \rightarrow \infty$ (achieved with $y_i \rightarrow \infty$) and similarly $\max_{y_i} g_i(y_i) \rightarrow \infty$ when $z_i < -\lambda_i$ (achieved with $y_i \rightarrow -\infty$). These cases are shown in Fig. 12. So, $F_2^*(z) = \infty$ for $z \notin [-\lambda, 0]$. For $z \in [-\lambda, 0]$ we compute the subdifferential of $g_i(y_i)$ (both shown in Fig. 12):

$$\partial g_i(y) = \begin{cases} z_i + \lambda & y_i < 1 \\ z_i + \lambda[0, 1] & y_i = 1 \\ z_i & y_i > 1 \end{cases}. \quad (36)$$

The optimality condition $0 \in \partial g_i(y_i^*)$ suggests that when $z_i \in (-\lambda_i, 0)$, $y_i^* = 1$. In other words, for that interval, we have $\max_{y_i} g_i(y_i) = z_i$. On the interval boundaries y^* is not unique (when $z_i = -\lambda$, $y^* \in (-\infty, 1]$, and when $z_i = 0$, $y^* \in [1, \infty)$), but $\max_{y_i} g_i(y_i) = z_i$ still holds. So the convex conjugate can be written as follows:

$$F_2^*(z) = \begin{cases} 1^T z & \text{if } z \in [-\lambda, 0] \\ +\infty & \text{else} \end{cases}. \quad (37)$$




Figure 12: Left: $g(y)$ plotted for $z \notin [-\lambda, 0]$. As noted in text, these functions are unbounded. Middle: $g(y)$ with $z \in (-\lambda, 0)$. This is a piecewise linear function with a unique maximizer (if $z \in \{-\lambda, 0\}$, the function is bounded, but the maximizer is not unique). Right: subdifferential of $g(y)$ on the left.

To solve for $\text{prox}_{\sigma F_2^*}(z)$, we first find $\arg\min_x \frac{1}{2\sigma} \|z - x\|_2^2 + 1^T x$ (ignoring the constraint $x \in [-\lambda, 0]$), and then project it onto the feasible set:

$$\text{prox}_{\sigma F_2^*}(z) = \arg\min_x \frac{1}{2\sigma} \|z - x\|_2^2 + F_2^*(x) \quad (38)$$

$$= \min(\max(z - \sigma 1, -\lambda), 0). \quad (39)$$

B.4. Layer unity prior and aggregated TV weights

In section 3.2 we made a claim that enforcing $\nabla c(x) = 0$ with a hinge loss penalty and associated cost $u(x)$ (nonzero for x where we want to enforce the constraint, and is zero elsewhere) is equivalent to increasing weights in TV regularization. We prove this claim here.

$\nabla c(x) = 0 \Leftrightarrow |\nabla c(x)| \leq 0$. Relaxed as a hinge loss penalty, this becomes

$$\int_D u(x) \max(0, |\nabla c(x)|) dx = \int_D u(x) |\nabla c(x)| dx. \quad (40)$$

But this is just TV regularization, and the objective already includes a penalty of the same form with weights $g(x)$. We conclude that the new “biased” penalty is $\int_D g'(x) |\nabla c(x)| dx$ with $g'(x) = g(x) + u(x)$.

C. Flow extrapolation via cross bilateral filter

In this section, we show additional examples of the result of flow extrapolation in the occluded region, which allows us to associate occluder points to occluded points. In each figure, the estimated motion field (v_t^{t+1}) and the extrapolated motion field (\hat{v}_t^{t+1}) are shown, below which regions of notable change are shown in individual in panels for closer inspection. In most cases, the extrapolation step locally improves optical flow.

For the car in Fig. 13, panels A-F show signs of improvement, where the flow in the occluded region becomes more similar to the “locally background” object. We note specific improvements as follows: A: the rear-view mirror is obtained. B: the motion in the region to the left of the car windshield becomes more similar to the background motion. C: incorrect optical flow in front of the car is fixed. D-E: most of the occluded region’s motion (in front of and below the car) becomes more similar to the background’s. F: the flow of the car wheel is independent of the vehicle motion, but is smoothed out. For the task of object segmentation at this scale, this is desired behaviour.

For the horse in Fig. 13, panels A-C show improvement. Box D is a failure example: extrapolated flow is worse than the original. Box E shows negligible improvement. Details follow: A: optical flow in the region around the head becomes similar to background. B: the erroneous motion in front of the horse’s nostril is removed. C: motion in front of the horse’s neck becomes is cleaned up. D: the left front leg of the horses is smoothed out, as the motion and appearance of are similar to background. When this occurs, the flow smoothes across the weak flow boundaries. E: the details on the tail improve slightly.




Figure 13: Two examples comparing the original optical flow field (v_t^{t+1}) (top left) and the resulting flow post-extrapolation (\hat{v}_t^{t+1}) (top right). Below each pair are panels highlighting regions where significant changes occurred. For each box, v_t^{t+1} is shown in the top row while the extrapolated result (\hat{v}_t^{t+1}) is shown below it.

C.1. Constraint perturbation

The ability to compute occluded-occluder pairs (y^c, y) reliably requires accurate optical flow, which may be difficult to obtain, especially in regions with large motion. Inaccurate flow may lead to an incorrect constraint (e.g. both points falling on the occluded region). To be robust against failure of optical flow, we assume that depth layers can be locally discriminated by their appearance, and locally perturb the constraints to ensure that both points fall on their respective sides of the occlusion boundary.

Specifically, we model the local appearance of local foreground (*occluder*) and local background (*occluded*) regions by a pair of Gaussian mixture models (GMM) with respective likelihoods f_F and f_B . Learning these models would be simple if we could obtain a set of samples from F and B regions. Since these are not known, we estimate f_F and f_B from nearby “occluder” and “occluded” points. We found that this approach works well in practice, despite the implicit assumption that the majority of these points fall into correct regions. To measure how well a point matches the appearance of *occluder* and *occluded* regions, we introduce likelihood ratios $\phi_F(x) \doteq \log \frac{f_F(I(x))}{f_B(I(x))}$ and $\phi_B(x) \doteq -\phi_F(x)$. If $\phi_F(x) > 0$ (resp. < 0), a point is likely in the foreground region (resp. in background region).

We would like to transform the pair (y^c, y) to ensure that y^c lies in foreground region and that y lies in background region, while penalizing large transformations. To achieve this, the procedure tractable, we restrict the transformation to be parameterized by translation, rotation, and uniform scale (i.e. a *similarity*). Formally, we can write down an optimization problem:

$$\max_g \phi_F(g \circ y^c) + \phi_B(g \circ y) - \|g\|, \text{ with } g \text{ a similarity transformation} \quad (41)$$

where we use the notation $g \circ x$ to denote a point x transformed by g . The first two terms measure how well the image intensities near transformed points match the appearance models. The deformation penalty is denoted by $\|g\|$. Once the best transformation g^* is found, we can use the transformed constraint pair ($g^* \circ y^c, g^* \circ y$). In practice, we compute the value of (41) for multiple local transformations, and choose the best one. An example of this procedure is shown in Fig. 14.

C.2. Local shape classifiers

In some image regions, poor motion estimation or excessive clutter can confuse the temporal integration and computation of segmentation weights ($g_t(x)$ in (11)) as both are based on pixelwise difference—this can make determining object boundaries difficult despite temporal integration. To reliably segment the objects from the background, we incorporate shape information over a local region near the object boundaries. In a fashion similar to [4], we employ a set of overlapping localized shape classifiers to help locally discriminate between foreground objects and the background. Each classifier learns an appearance




Figure 14: *Constraint perturbation helps clean up poorly estimated constraints. Left: original constraints before perturbation. Middle: improved constraints. Right: several occluder-occluded pixel pairs where perturbation ensures that estimate of the occluder’s position lies on the foreground object (the local foreground probability map is shown under ϕ_F).*

model and a shape model of the object within a small region, and then propagates this information to the next frame by locally adjusting strength of the foreground prior κ_t .

After the first frame is partitioned, a set of local classifiers is instantiated along object contours to track the boundary of the foreground object o_{fg} within a small window W (31 x 31 pixels). Local shape is determined by the mask of o_{fg} in W , denoted $m(x)$, and the local appearance of o_{fg} is modeled by a GMM (3 modes). A confidence c_{fg} on how well the color model discriminates between o_{fg} and background is computed as

$$c_{fg} = 1 - \frac{\int_W |m(x) - p_c(x)| w_c(x) dx}{\int_W w_c(x) dx}, \quad (42)$$

where p_c is the foreground probability computed from the appearance GMM and $w_c(x) = \exp(-d^2(x)/\sigma_c^2)$, which weights the contribution of each pixel based on $d(x)$ (the distance between x and the foreground-background boundary, computed using a distance transform). σ_c is set to half the window size (for us, 15 pixels) (see (2) in [4] for further details). Next, the local classifier is warped forward by the motion of o_{fg} . Based on c_{fg} , appearance and shape models adaptively combined to provide a prior on which pixels in W are likely to be foreground

$$p_{fg}(x) = c_{fg} p_c(x) + (1 - c_{fg}) L(w_t^{t+1}(x)), \quad (43)$$

where $L(w_t^{t+1}(x))$ is the binary mask $L(x)$ warped into the current frame. Finally, the contributions from each local classifier are combined together and added to $\kappa_t(x)$ in (11).

After the first frame, each classifier is propagated forward until its support contains no foreground objects, at which point it is dropped. New classifiers are instantiated where object boundaries in the current frame’s segmentation are not covered by any local classifiers. See [4] for further details. Generally, we find that these classifiers improve the segmentation of foreground objects in regions where a strong boundary is not discernible as shown in Fig. 15. In addition, for small objects moving quickly, these classifiers help preserve the object as the foreground prior can sometimes be eroded away as shown in Fig. 16.




Figure 15: Two examples of the local shape classifiers boosting the foreground prior (κ_t) near object edges, where κ_t is computed without the classifiers on the left and incorporating them (with the classifier windows overlayed in light gray) on the right. Notice how κ_t better captures the shape of the woman's hat behind the bush (left) and the legs of the horse (right).




Figure 16: Examples of local shape classifiers preserving small objects when the foreground prior term alone would erode away. From left to right: κ_t without local classifiers, the corresponding layer segmentation, κ_t incorporating the classifiers (with the windows drawn in light gray), and the resulting layers. The camel's head and legs are better captured with the help of local shape classifiers (row 1). Without them, much of the object may go missing (row 2) or even become lost altogether (rows 3–4). These local shape classifiers are beneficial to long-term temporal consistency.